Chapter 5
Why the Child Holded the Baby Rabbits: A Case Study in Language Acquisition

Steven Pinker

5.1 Introduction: The Creativity of Language Users

Human language is one of the wonders of the natural world. Unlike other species' calls and cries, which convey a fixed set of messages—such as warnings of danger or claims to territory—the noises that come out of our mouths can convey an unlimited number of different, precise, structured, brand-new propositions. I can put together sentences that can tell you anything from how to build a small thermonuclear device in your basement, to the mating habits of the octopus, to the latest twists in your favorite soap opera plot. You would never have heard these sentences before, but you would recognize them as English and understand their meanings. The number of sentences that a human can produce and understand is astonishing—a hundred million trillion different sentences twenty words or less, according to one estimate. Indeed, a person is capable, in principle, of producing an *infinite* number of sentences (putting aside the fact all humans are mortal). By the same logic that allows us to say that there are an infinite number of numbers (if you ever think you have listed them all, I can create a new one by adding 1 to the largest), there are an infinite number of sentences: if you ever think you have listed them all, I can create a new one by adding *He wrote that* ... to the longest.

A human head is not big enough to store an infinite number of sentences, or even a hundred million trillion sentences. So what we know when we know a language is a program, or recipe, or set of rules, that can string words together in an unlimited number of systematic combinations. This set of rules is sometimes called a *mental grammar* (which is not to be confused with the prescriptive, or school, grammars that dictate how one "ought" to write and talk in standard English, a different matter entirely). A mental grammar contains rules that define sentences by concatenating

Preparation of the chapter was supported by NIH grant HD 18381, NSF Grant BNS 91-09766, and the McDonnell-Pew Center for Cognitive Neuroscience at MIT.

pieces of sentences, and define those pieces in terms of smaller pieces, and so on. For example, a sentence consists of a subject and a predicate, and a predicate consists of a verb and an object and possibly another whole sentence embedded inside—a trick called "recursion" that can create arbitrarily long sentences with sentences inside them, such as *I think that he thinks that she thinks that....* The nature of mental grammars is discussed in chapter 10; how they are put to use is discussed in chapter 8. In this chapter we will explore a crucial aspect of mental grammars: how children develop them.

The way that babies develop language is as wondrous a phenomenon as the way that adults use language. Indeed, the linguist Noam Chomsky revolutionized the study of language in the late 1950s by pointing out that language acquisition is the key to understanding the nature of language.

Obviously, children are not born with the mental rules for any particular language—a baby will acquire English or Japanese or Cherokee or Swahili, depending on where it is brought up; but, just as obviously, babies are not taught the rules of those languages. Parents just talk to their children; they do not give them grammar lessons. Somehow, the children must distill the rules out of their parents' sentences. And this is a difficult problem for any computational system to solve. By the time they are in their threes, normal children speak their language fluently. They have heard only a finite number of sentences from their parents—however many they can listen to in those three years—but have developed the ability to produce and understand an infinite number of new sentences. They cannot just memorize their parents' sentences; they have to take leaps, and make guesses about what other kinds of word combinations are possible in the language. Logically speaking, there are an infinite number of legitimate ways of making such leaps. A child *could* conclude that English consists of all the sentences heard so far, plus their right-to-left mirror images (perhaps to convey the opposite of the sentence's original meanings), or of all the sentences heard so far, plus the same sentences with their auxiliaries omitted (*He taken a ride*), or of all those sentences, plus the same ones with the word *what* stuck in front of them (to turn them into questions), and so on. All those leaps but one—the correct mental grammar for English—are wrong, but they are all consistent with the parents' speech. How do children make just the right leap?

Chomsky pointed out that the only way for children to solve this problem is to be innately equipped with mental machinery that forces them to analyze sentences only in certain ways, and thus to pull out only certain kinds of rules and principles—the ones that in fact underlie humanly possible languages. He called this machinery *universal grammar.* Among the psychologists and linguists who study language, there is a

controversy over whether universal grammar is an independent "mental organ," or a special case of more general principles of mental life (used in other domains such as problem solving and social interactions); but it is a logical necessity that children's mental learning mechanisms be constrained in *some* way, for otherwise they could not generalize correctly beyond their parents' sentences to the rest of the language.

5.2 Overregularization: A Case Study of Grammatical Creativity

In this chapter we will try to understand the process of language acquisition by focusing on a very clear example where we can catch children in the act of generalizing beyond what they hear from their parents. By their late twos, all normal children produce sentences like the following:

> My teacher holded the baby rabbits and we patted them.
> Hey, Horton heared a Who.
> I finded Renée.
> I love cut-upped egg.
> Once upon a time a alligator was eating a dinosaur and the dinosaur was eating the alligator and the dinosaur was eaten by the alligator and the alligator goed kerplunk.

Obviously, no adult who speaks standard English goes around saying *heared, finded,* or *goed,* so children could not have memorized these verbs from their parents. They must have abstracted a rule—something like "to form the past tense, add the *-ed* sound to the end of the verb"—and are trying it out on "irregular" verbs that don't ordinarily submit to the rule, like *to go.* (They also try it out on verbs that they have created themselves, like to *cut-up,* a misanalysis of the verb-plus-particle combination *cut + up.*) Indeed, experimental psychologists have shown, in a different way, that preschool children are quite willing to apply the past tense rule to new verbs to create brand-new past tense forms. In 1958 Jean Berko (now Berko-Gleason) ran a simple but ingenious experiment. Showing children a picture of a man swinging a club, she said, "Here is a man who knows how to *rick.* He did the same thing yesterday. He...." Most of the preschool children eagerly filled in the blank with the correct new form *ricked.* Similarly, when asked to complete a sentence about more than one *wug* (a little cartoon bird), they created *wugs;* and in honor of that bird the experimental paradigm is now called the *wug*-test.

5.2.1 The Course of Rule Development

Children's willingness to apply the English past tense rule to inappropriate and novel verbs offers us a window on their acquisition of productive

grammatical rules. We can even see the acquisition of the rule happen before our eyes, as we follow children's early language development. Children produce isolated words around the age of one and combine them into two-word microsentences like *See baby* and *More milk* around age one and a half. By their twos, children are producing longer and more complex sentences, and they are beginning to supply the obligatory functional morphemes like *-ing*, *-ed*, and *-s*, and auxiliaries. Somewhere between the ages of late one and late two, the first "overregularization error" (as forms like *goed* and *holded* are called) appears—a clear sign that the child has acquired something like the past tense rule.

Interestingly, children do not avoid past tense forms before using the rule. Indeed, for months on end they can use irregular past tense forms correctly, like *held*, *went*, and *heard*, together with a smattering of regular forms, like *played* and *used*. Presumably, they simply memorized those forms from their parents' speech and were using them just like any other words (with the "pastness" just a part of the words' meanings). At some point they "notice" (not consciously, of course) that many words come in two, ever-so-slightly different versions: *walk* and *walked*, *use* and *used*, *play* and *played*, *push* and *pushed*, and so on. Of course, the child could, logically speaking, simply chalk these up to variations in pronunciation or speaking style. But if universal grammar says something like "There exist rules for adding affixes onto words to signal tense, number, person, gender, and so on—Watch out for them!" the child would therefore be impelled to look for some systematic rule underlying the variation. By subtracting *walk* from *walked*, *push* from *pushed*, and so on, the child can pull out the *-ed*; by checking for what correlates with the use of the *-ed* versions, the child can notice that Mom and Dad always use it when they are describing events that are over and done with, and therefore that adding *-ed* signifies "past tense." (The child must also be forced to assume that the rule is *obligatory*—you can't say *I eat breakfast earlier this morning*, even though the meaning is completely clear.) Having acquired the rule, children can feed any verb into it, regular or irregular; thus, they start saying *goed* and *holded* in circumstances where earlier they might have said *went* and *held*.

Child psychologists call the whole sequence "U-shaped development" because if you plot the child's performance over time—specifically, the number of past tense forms of irregular verbs that are correct past tense forms, as opposed to overgeneralizations—the curve starts at 100 percent and then dips down (in one sense, the child is getting *worse* while growing older) before rising back to 100 percent by adulthood. U-shaped development is always an interesting phenomenon because it reflects some qualitative change in the child's development, not just better and better approximations of adult performance.

5.2.2. Explaining Rule Development Is Not So Easy

The story I just told is an oversimplification in many little ways that we can ignore for the time being (for example, the past tense suffix is actually pronounced in three different ways in *walked, jogged,* and *patted,* and the child must cope with this), and in one big way that we cannot ignore. Here is the problem. I have told you that children start to say *holded* because they have acquired the "add *-ed*" rule. But adults have the "add *-ed*" rule, too, but they *don't* say *holded.* If they did, we would not call the child's form an "error" to begin with, and I could never have used these forms as an example of children's creative use of grammatical rules, the whole point of the example! Something very important is missing from the story: the difference between children and adults, and how children eventually eliminate that difference as they grow up.

A first guess might be that language development is driven by communication: children improve their language in directions that allow them to communicate their wishes more effectively. Wrong! There is nothing unclear about the meaning of *holded.* In fact, as long as children are willing to make overregularization errors, their language is *more* communicative than adults'. There are about twenty-five irregular verbs in adult English that don't change their forms in the past tense, like *cut, set,* and *put;* thus they are ambiguous between past and nonpast. *On Wednesday I cut the grass* could mean last Wednesday, next Wednesday, or every Wednesday. The childlike form *On Wednesday I cutted the grass* could mean only a preceding Wednesday. A language is obviously a very powerful form of communication, but children must not be acquiring all of its details *because* they help in communication; they learn adult English because they can't help it, and just use it, whatever its strengths and weaknesses.

A second guess is that adults don't say *holded* and *goed* because they have never heard other adults say them. Wrong again. This hypothesis predicts that adults could not pass a *wug*-test, which is absurd. New verbs enter the language all the time, like *to diss, to mosh, to snarf, to frob, to wild, to mung, to flame,* and so on. Surely adults who learn *diss* in the present tense do not have to wait to hear someone else say *dissed* before they can say it; the forms *dissed* and *dissing* are created automatically. So, we cannot explain why adults avoid *holded* by saying that they've never heard anyone else say *holded.* Adults have never heard anyone else say *dissed,* either, but they don't avoid *dissed.*

The problem with overregularizations is not that they have never been heard before; it's that the irregular counterpart *has* been heard. There must be a piece of adult psychology that causes the experience of hearing an irregular form like *held* to block the subsequent application of the regular "add *-ed*" rule to that item. Some linguists have called this mechanism the

Blocking principle. This means that an idiosyncratic form listed in the mental dictionary as corresponding to a particular grammatical modification of a word (past tense, in this case) blocks the application of a general rule that would effect the same grammatical modification. Thus, *held,* listed as an idiosyncratic past tense form of *hold,* blocks application of the regular past tense rule, preempting *holded; geese,* listed as a plural of *goose,* blocks *gooses; better,* listed as a comparative of the adjective *good,* blocks *gooder.*

So, maybe children lack the Blocking principle and have to learn it. This explanation is no help, though, until we figure out how they *could* learn it, and that raises a very big problem related to the logic of learning. To learn the Blocking principle, children have to know that forms like *holded* are ungrammatical. Now, they can't conclude that *holded* is ungrammatical simply because they have not heard adults use *holded;* they have not heard adults use *ricked* either, and they are happy to say *ricked* (and must continue to make such leaps as adults in order to be able to produce *flamed, dissed, moshed,* and so on. The only way of knowing that *holded* is ungrammatical is to use it and be corrected, or to get some other negative feedback signal from adults like disapproval, a puzzled look, or a non sequitur response. But it does not appear that parents provide such systematic feedback (which is sometimes called *negative evidence*). After all, if they did, they would be correcting or disapproving their young children almost all the time, since most of toddlers' sentences are ungrammatical in some way or another. No, parents care much more about truth and good behavior than grammaticality, as we can see in the following typical dialogues:

Child: Mamma isn't boy, he a girl.
Mother: That's right.

Child: And Walt Disney comes on Tuesday.
Mother: No, he does not.

The psychologists Roger Brown and Camille Hanlon (1970) verified this impression quantitatively. They examined transcripts of the natural conversations of three children they were studying, with the pseudonyms Adam, Eve, and Sarah (the first three names in the Bible), and divided their sentences into grammatical and ungrammatical lists. For each sentence they checked whether the parent expressed approval (such as "Yes, that's good") or disapproval. The proportion was the same for grammatical sentences and ungrammatical ones, so the child can get no information about grammar from such responses. They also checked whether children might learn about the state of their grammar by noticing whether they are being understood. They looked at children's well-formed and badly

formed questions, and whether parents seemed to answer them appropriately, as if they understood them, or with non sequiturs. Again, there was no correlation; *What you can do?*, like *holded*, may not be English, but it is perfectly understandable. We can conclude that children do not seem to have access to information about which of their words or sentences are ungrammatical.

To be fair, there has been some controversy about this point, and a few psychologists have tried to show that there are very subtle and probabilistic patterns in some parents' reactions to their children that might give the children information about which of their utterances are ungrammatical. But the psychologist Gary Marcus (1993) has shown that even these patterns are so unreliable as to be useless (children would have to repeat each error verbatim hundreds of times to get reliable feedback information, because most of the time parents react the same way to grammatical sentences as to ungrammatical ones). The physician-psychologist Karin Stromswold (1995) has an even more dramatic demonstration that parental feedback cannot be very important. She studied a child who, for unknown neurological reasons, was congenitally unable to talk. He was an avid listener, though, and when tested was able to understand complicated sentences perfectly and to judge accurately whether a sentence was grammatical or ungrammatical (including correct dismissals of many overregularized test items like *holded*). The boy's abilities show that children certainly do not *need* negative evidence to learn grammatical rules properly, even in the unlikely event that their parents provided it.

5.3 A Simple Explanation of Overregularization—That Works

Logically speaking, given the knowledge that forms like *holded* are not English, children could learn the Blocking principle. Since this is unlikely—children have no direct evidence that *holded* is not English—we can turn to an alternative: Given the Blocking principle, children could learn that forms such as *holded* are not English; they would just have to hear their parents say *held*. Perhaps Blocking is part of Chomsky's innate universal grammar and is a cause, rather than an effect, of the child's language learning.

But we need one more assumption to get this theory to work. If children had a Blocking mechanism, they should never say *holded* to begin with—having heard their parents say *held* even once, they should block the rule from applying to it! The extra needed assumption comes from an uncontroversial principle of the psychology of memory, known for over a hundred years. People do not remember an arbitrary pairing (like a name with a face, or a treaty with a date) perfectly on a single exposure. It often

takes repeated encounters, with the probability of successful retrieval increasing with each encounter (presumably reflecting an increase in the "strength" or clarity of the trace of the pairing as it is stored in memory). Now children, by definition, have not lived as long as adults. So children have experienced everything in life fewer times than an adult has, including hearing the past tense forms of irregular verbs. If children have heard *held* less often, their memory trace for it will be weaker, and retrieval less reliable. Some of the time, when they are trying to express the concept of holding in the past, *held* will not pop into mind (or at least, not quickly enough to get it out in that sentence). If they are at an age at which they have already acquired the regular past tense rule, they will apply it to *hold*, creating *holded*, so as to satisfy the constraint that tense must be marked in all sentences. Prior to that age, when they failed to retrieve *held*, they had no choice but to say *hold*.

This account, which combines a simple idea from linguistics (Blocking as part of the child's innate universal grammar) and a simple idea from psychology (memory retrieval is fallible, with a probability that increases with increasing number of exposures to the pair), has a number of advantages. For one thing, it removes the paradox about why children get worse, and how they get better in the absence of parental feedback. The irregular forms that the child uses early on do not go anywhere once the child has acquired the regular rule, nor are those forms ever incapable of blocking overregularization: they just have to be retrieved from memory to be able to do the blocking, and they are retrieved probabilistically. The cure for overregularization is living longer, hearing the irregulars more often, and consolidating them in memory, improving retrievability.

Indeed, according to this account, we do not even need to posit any qualitative difference between the mind of the child and the mind of the adult, and so the account has parsimony on its side. In fact, it is deducible from the very logic of irregularity, supplemented only by the simple assumption that memory is fallible. What is the past tense form of the verb *to shend*, meaning "to shame"? If you answered *shended*, then you have overregularized; the correct form is *shent*. Of course, this "error" is not surprising. Irregular forms are not predictable (that is what "irregular" means), so the only way you could have produced *shent* was if you had previously heard it and remembered it. But you have heard it zero times, so you can't have remembered it. Now, if in two years you were asked the question again and overregularized it once more, it would still not be surprising, because you would have heard it only once. Many verbs for a given child will be like *shent* for you: never heard; heard but not attended to; heard and attended to, but not enough times to be able to recall it on demand reliably. The *holded* phenomenon has been demystified from beginning to end.

5.4 Evidence for the Blocking-Plus-Retrieval-Failure Theory

Even the most beautiful theory can be killed by an ugly fact, so we have to examine the facts before we can be content that we have understood this example of language acquisition in action. Gary Marcus and I have assembled ten facts that we think support the overall story fairly nicely (Marcus et al. 1992).

First, we looked at the actual *rate* of making the errors. The theory predicts that the child's linguistic system is at all times designed to suppress regularization of verbs remembered to be irregular. This suppression cannot be perfect because the child's memory is not perfect, but it is as good as the child's memory retrieval process. If we assume that children's memory for words, though imperfect, is good (the child is, after all, successfully using thousands of words, and acquiring them at a rate of approximately one every two waking hours), then overregularization should be the exception, not the rule, representing the occasional breakdown of a system that is built to suppress the error.

Marcus and I examined computer-based transcripts of the spontaneous speech of twenty-five individual children, together with transcripts of another fifty-eight children speaking with one another in various-sized groups. We looked at their rates of overregularization, calculated as the number of errors (such as *holded* and *helded*) divided by the number of opportunities for the child to make an error, which we conservatively estimated by adding the number of errors to the number of correct past tense forms like *held*. The mean rate across children was only 4.2 percent; the median only 2.5 percent! That is, more than 95 percent of the time when a child uttered the past tense form of an irregular verb, it was the correct form like *held*, not an error like *holded*.

Moreover, once children began to overregularize at all, they kept doing it at this low rate pretty steadily throughout the preschool years (2 to 5). That is, the 2.5 percent rate is not a misleading statistic resulting from averaging together, say, thirty-nine months in which the child was 100 percent correct and one month in which the child had completely reorganized his or her grammar and was 100 percent incorrect. The top curve in figure 5.1 shows the percent correct (excluding ambiguous tenseless forms like *hold*) for the boy called Adam. You can see the transition from correct performance to occasional errors (the left arm of the very roughly "U-shaped" development); but once the errors are made, they are made at roughly the same rate for years. This is just what we would expect if the errors were sporadic malfunctions, not a product of a qualitatively different kind of system.

The low rate is also not an artifact of averaging together a few verbs that are always overregularized (which would be a mystery on our simple

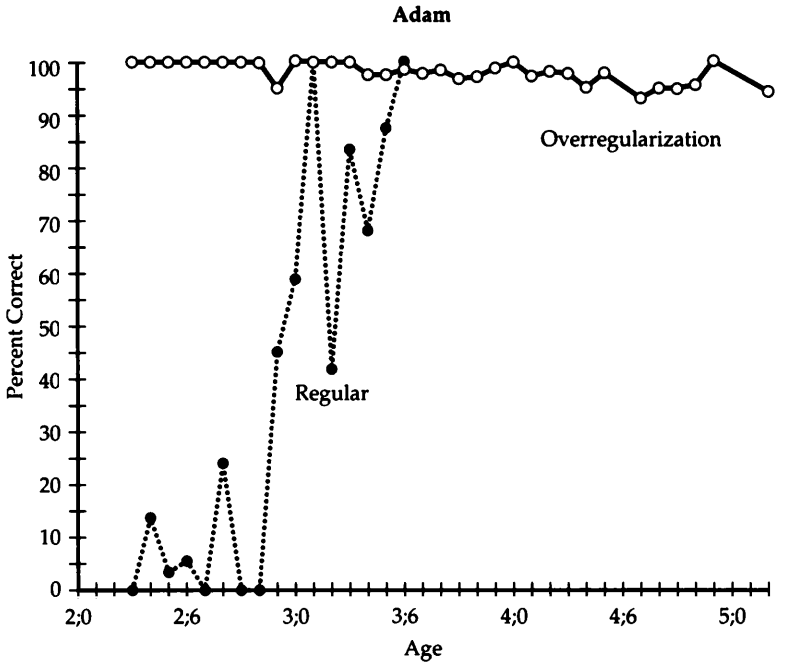

Figure 5.1
"Overregularization" curve: Percentage of Adam's irregular past tense forms (e.g., *held* or *holded*) that are correct (e.g., *held*). Equivalent to 100 percent — overregularization rate.
"Regular" curve: Proportion of Adam's regular verbs in obligatory past tense contexts (e.g., *He walk* or *He walked*) that were correctly marked for tense (e.g., *He walked*).

theory) with many verbs that are never overregularized. We found that no verb was immune to overregularization, not even the verbs that a child had used correctly before beginning to overregularize: a child might use *felt* when young, then both *felt* and *feeled* when slightly older. Nor did we find verbs that one can conclude are *always* overregularized. There were a few, to be sure, but they were always verbs represented by tiny, possibly misleading samples. A child might have said *holded* once and *held* zero times, for an error rate of 100 percent; but, if the tape recorder had happened to catch him on a different day, it could have been *held* once and *holded* zero times, for an error rate of 0 percent. Whenever a child used a verb in the past tense often enough for us to get a reliable sample, avoiding this problem, we virtually always found that the error percentage was in the single digits. Once again, it looks as though overregularization is fairly haphazard from one moment to another. In fact, children can use the correct and overregularized version of the same verb in quick succes-

sion, like this: "Daddy comed and said 'hey, what are you doing laying down?' And then a doctor came . . ."

Second, there is some predictability as to which verbs are overregularized more often, and it fits nicely with the theory that memory retrieval failure is a prime culprit. We counted how often the children's parents used each irregular verb in the past tense (fortunately, the parents' speech, not just the children's, was included in the transcripts). The prediction is: if a parent uses *brought* and *told* more often than, say, *froze* and *wound*, the child should have a stronger memory trace for *brought* and *told* than for *froze* and *wound*, and hence should say *bringed* and *telled* less often than *freezed* and *winded*. Actually, we compared ninety irregular verbs, not just four, and found that, as predicted, every single child made more errors on the verbs that the parents used less often in the past tense: the correlations between overregularization rates and frequency in parental speech averaged −.33.

A third kind of evidence comes from examining what happens to the *regular* verbs at the point at which the child first makes an error with the irregulars. The lower curve in figure 5.1 plots the percentage of time Adam correctly marked regular verbs for past tense (that is, with *-ed*) in contexts that required past tense marking: that is, how often he said *Yesterday he walked* as opposed to *Yesterday he walk*. As you can see, before beginning to overregularize, Adam left regular verbs unmarked most of the times in which they should have been marked. After he began overregularizing, he started to mark them more often than not. Indeed, the curve for correct marking of regular verbs ascends steeply at just the point at which Adam first overregularized. If we assume that this is the age at which he acquired a well-functioning "add *-ed*" rule, the data can be explained nicely. In the first phase Adam was stuck with the past tense forms that he had memorized, both regular and irregular. If he lacked the past tense form for a verb, or if he had a past tense form in memory but could not recall it instantly, then that was just too bad for the syntax of the sentence; he had no choice but to use the bare verb stem. But in the second phase he had the *-ed* rule available to fill any memory vacuum. This allowed him to mark regular verbs correctly *and* to mark irregular verbs incorrectly, as overregularizations—as if marking an irregular incorrectly is better than not marking it at all. The graph, in a sense, catches Adam in the act of growing a rule.

A fourth kind of evidence confirms that children are not radically reorganizing their grammars when they begin to overregularize (which would by mysterious) but are simply adding to it. Remember that at the ages in which children overregularize, they do so only around 4 percent of the time. Where do those four percentage points' worth of errors come from? Are they produced in the kinds of sentences for which Adam, before he acquired the rule, would have produced a correct irregular form like *held*?

Or are they produced in the kinds of sentences for which he would have produced a bare stem like *hold*? That is, are the errors driving out correct forms like *held* (a mysterious step backward on the road to adult language), or are they merely driving out a different kind of error, *hold*? We cannot know with certainty, but here is a statistic suggesting that these new errors are taking the place of old errors. In phase 1, Adam used correct forms like *He held it* 74 percent of the time, and made errors like *He hold it* 26 percent of the time. In phase 2 he used *He held it* 89 percent of the time, *He hold it* 9 percent of the time, and *He holded it* 2 percent of the time. It looks as if the overregularizations are occurring at the expense of the forms that used to be bare-stem errors, not at the expense of correct forms; the correct forms just get likelier and likelier as the child gets older. We thus see no real regression or reorganization or backsliding, just a sophisticated kind of error (since it represents an effort to mark tense) replacing a simpler kind (no marking at all).

Fifth, we looked at the results of a certain kind of experiment, in which children are asked to act like little linguists and tell the experimenter whether certain words are grammatical or ungrammatical. Of course, children do not know what the word *grammatical* means, so clever games have to be invented to tap their intuitions of grammaticality indirectly. Typically, the experimenter has a puppet say various things, and the child is asked to judge whether the puppet "said something silly."

Now, one has to be careful in interpreting any experiment in psychology in which the subject gives the equivalent of a "yes" or "no" answer in trying to detect the presence of some condition (perception researchers call these "signal detection" problems). In most cases the subject will have some degree of confidence in whether the condition is present (say, whether a light came on, or whether a sentence is grammatical), and then will have to decide whether that confidence level is high enough to warrant saying "yes" or "no" (or, for the child, "not silly" or "silly"). That decision will be influenced not only by the subject's actual degree of certainty but also by the subject's willingness to say "yes" or "no" more than a certain proportion of the time when there is any doubt at all. This will depend on the subject's feelings of the costs and benefits of saying "yes" when the condition is absent (a false alarm) versus saying "no" when the condition is present (a miss). In these experiments in particular, each child will have a different willingness to say "silly!" depending on his sense of politeness, fun, and many other factors that we cannot precisely measure. The point of this methodological digression is that it is not easy to interpret the percent-correct measure from any psychology experiment in which the subject gives a yes-or-no answer. It is generally sounder to *compare* the subject's "yes" rate when the condition is actually present with the "yes" rate when the condition is absent.

In these experiments (conducted by the psychologist Stan Kuczaj, 1978), the children, who ranged in age from 3 to 9, failed to say "silly" to many overregularizations—as many as 89 percent for some children. Does this mean that their mental grammars actually accept overregularizations? Not necessarily—they could just be shy about calling the experimenter's puppet "silly." A better way to look at the data is to compare how often they said "silly" when the puppet used an overregularization and when the puppet used a correct past tense form. In virtually every test the children were statistically *more* likely to call the overregularized form "silly"; this suggests that their grammars really did deem those forms worse, despite their occasional errors.

A sixth and related kind of evidence is anecdotal. Some psycholinguists have tried using their own children's overregularizations when talking to the children, just for fun. The children were not amused:

Parent: Where's Mommy?
Child: Mommy goed to the store.
Parent: Mommy goed to the store?
Child: NO! (*annoyed*) Daddy, *I* say it that way, not you.

Child (a different one): You readed some of it too . . . she readed all the rest.
Parent: She read the whole thing to you, huh?
Child: Nu-uh, you read some.
Parent: Oh, that's right, yeah. I readed the beginning of it.
Child: Readed? (*annoyed surprise*) Read! (*pronounced "red"*)
Parent: Oh, yeah, read.
Child: Will you stop that, Papa?

Again, it looks as though the children, at some level in their minds, compute that overregularizations are ungrammatical even if they sometimes use them themselves.

A seventh bit of evidence also comes from adult overregularizations—but this time accidental ones. If you listen long enough, you will hear *adults* saying things like *drived* and *blowed* as slips of the tongue. (Occasionally I hear my students chuckling softly in lectures, and one will later point out to me that I have done it myself.) Not very often, of course—perhaps a few thousandths of a percent of the time—but the fact that adults make them at all suggests that the errors themselves do not have to be the product of a different kind of mental grammar. Indeed, the psychologist Michael Ullman (Ullman et al. 1993) has found that as adults get older and their memory gets less and less reliable, they begin to make overregularizations more often, at least in experiments. (So perhaps the full developmental curve is tilde-shaped, not U-shaped.)

An eighth kind of evidence involves a gray area where adults, like children, use both overregularized and correct forms of a verb, but even the adults are not sure which is "correct." Most Americans are unsure about the best past tense form of certain irregular verbs. For example, is it *dreamed* or *dreamt*? *Dived* or *dove*? *Leapt* or *leaped*? *Slitted* or *slit*? *Slayed* or *slew*? *Strided* or *strode*? These indecisive verbs are lower in frequency of use (at least in the written language) than verbs for which there is no doubt among speakers (such as *went, came, told,* and the like; Marcus et al. 1992; Ullman 1993). And in a study asking people to give ratings to the regular and irregular forms of several hundred verbs (Ullman 1993), we found that the more often the irregular version of one of these verbs is found in the written language, compared with the regular version, the better the irregular form sounded to the people.

The explanation is as follows. Say you have heard *strode,* an uncommon verb, only a few times in your life—more often than *shent,* of course, but far far less often than *held.* You would have a weak, fuzzy memory trace for *strode.* You would recognize it when you heard it, but when you tried to say it, you'd have a little voice in your mind's ear shouting "*strode!*" but not so strong as to block the regular rule from applying to *stride;* you might thus go ahead and say *strided.* For verbs uncommon enough for many people to be in doubt about the irregular, but still lively enough for the irregular form to be recognizable, this can sometimes lead to a chaotic situation in the language community as a whole. Some people use *strided,* others use *strode,* and still others—hearing various parents and neighbors using one or another without rhyme or reason—simply memorize them both and use them interchangeably in their own speech.

But for even rarer verbs, adults' systematic "overregularizations" could create a vicious circle: the adults would use the irregular form less and less; thus their neighbors and children would themselves hear the irregular form used less and less, causing their memory traces for the irregular to be weak, causing *them* to use the irregular form less (overregularize it more), causing *their* neighbors and children to hear it less, and so on. An irregular form that falls below a certain level of usage could simply disappear from the language after a few generations. This is precisely what has happened in the history of English. Old English (the version of English spoken before the twelfth century) had many more irregular verbs than modern English does, such as *chide-chid, abide-abode, cleave-clove,* and *geld-gelt* (look it up). As time passes, some verbs wane in popularity, and one can imagine a time when, say, the verb *to geld* had slipped so far that a majority of adults could have lived their lives seldom having heard its past tense form *gelt.* When pressed, they would have used *gelded;* the verb had become regular for them—and for all subsequent generations.

The linguist Joan Bybee has tested this scenario using "diachronic" evidence (lingo for historical change, as opposed to the "synchronic" evidence about the state of the language at any one time), and this gives us a ninth piece of evidence. Bybee examined thirty-three surviving verbs that were irregular in Old English. Fifteen have come through in Modern English as irregular verbs; eighteen have become regular (such as *chide*). The surviving irregulars have a mean frequency in modern written English of 515 uses per million words of text (137/million for past tense forms alone); the regularized verbs have a mean frequency of 21/million (5/million in the past tense). That is, the English language over the centuries behaves just as children and adults do today: it regularizes the rarer irregular verbs. Presumably, this is because children today, adults today, children of yesteryear, and adults of yesteryear all have the same psychology of language.

The tenth kind of evidence comes form the statistics of twentieth-century English vocabulary. If the survival of an irregular form against the pressures of children's (and adults') regularizing them depends on high frequency, then we should find that clear irregular verbs tend to be the commonly used ones in the language. Regular verbs, on the other hand, can be formed by a rule whenever a past tense form is needed, and hence a regular verb can have a past tense form regardless of its frequency; it can be common or rare. (Note that it is not a logical necessity that every word have a predictable inflected form. Often a language will simply lack a rule for some category, as English does for inhabitants of a city. While everyone knows that residents of London are *Londoners*, residents of Boston are *Bostonians*, not *Bostoners*, and non-natives generally shrug when asked for the names of residents of St. Louis, Falmouth, Tel Aviv, and so on. Verbs are different: no matter how obscure, everyone can guess a past tense form for any verb.)

Computational linguists at Brown University have assembled the frequency statistics for one million words of English text, including newspapers, popular books, textbooks, and many other sources. The texts contained 4,291 different verbs. Here are the top ten in frequency:

1.	be	39,175/million
2.	have	12,458
3.	do	4,367
4.	say	2,765
5.	make	2,312
6.	go	1,844
7.	take	1,575
8.	come	1,561
9.	see	1,513
10.	get	1,486

Note that all ten are irregular! (Just as striking is the fact that the third person singular present inflection, whose regular form is *-s* as in *He walks*, has only four irregular exceptions: *is, has, does, says*. They constitute the four most frequent verbs in English.) There are a thousand verbs tied for *least* frequent in the Brown sample, all having a frequency of one occurrence per million. Here are the first few:

3791. abate	1/million
3791. abbreviate	1
3791. abhor	1
3791. ablate	1
3791. abridge	1
3791. abrogate	1
3791. acclimatize	1
3791. acculturate	1
3791. admix	1
3791. adsorb	1

As you can see, these are all regular. More precisely, 982 of the bottom 1,000 are regular, and one is a genuine irregular—the slightly quaint-sounding *smite-smote*. Another sixteen irregulars snuck in (a nice colloquial irregular!) in disguise: two-piece words that consist of an irregular *root* (the smallest unanalyzable piece of a word) plus a prefix: *bethink, forswear, inbreed, mis-read, outdraw, outfight, overbear, overdrive, overlie, overwrite, pre-sell, regrind, spellbind, unbend, unbind, unwind*. Incidentally, some of these prefix-plus-root combinations show some interesting properties of the human language system. The regular form sounds completely awful (*bethinked, forsweared, inbreeded, mis-readed*, and so on), despite the low frequency of the whole irregular form. The fact that the *roots* are high in frequency (*think-thought, swear-swore*, and so on) seems to be enough to block the regular rule. This shows that people mentally analyze words into a prefix plus a root, when they can; it is the root that does the blocking (so a familiar root inside an unfamiliar word, as in *bethink*, blocks regular *bethinked*). But for many of these verbs, the irregular form also sounds strange, like *bethought, forswore, overdrove, spellbound*, presumably because of the unfamiliarity of the entire two-piece word. Because we snip words into their parts when figuring out how to inflect them, these verbs have no perfectly natural sounding past tense forms: the regular is blocked by the fairly frequent irregular root, but the irregular is weakened by the unfamiliarity of the word as a whole.

5.5 Past Tense Overregularization and Connectionist Modeling

I have been writing as if children's creative errors with *-ed* can be accounted for only by assuming that there is an "add *-ed*" rule in the child's

head. But there are alternatives. Perhaps children use some kind of *analogy* from words they already know. They might say *holded* because *hold* sounds like *fold, mould,* and *scold,* and the past tenses of those verbs are regular *folded, moulded,* and *scolded.* Of course, some verbs, like *sing* and *ring,* do not rhyme exactly with any common regular verbs, but children still might be using bits and pieces of partially similar verbs they have heard, like *sipped, banged, rimmed,* and *rigged,* to assemble the analogous *singed* and *ringed.*

Until recently, no one knew how exactly this kind of complex analogizing would work, so it was not clear what to make of this alternative. But the school of computer modeling called *parallel distributed processing* (PDP), *connectionism,* or *artificial neural networks* (see the chapter by James Anderson in volume 4) introduced a new set of possible mental mechanisms to the field of psychology. One of them, called a *pattern associator,* is designed both to memorize specific pairs of items in a training set *and* to generalize them to new forms using complex partial analogies. In such networks an input is usually represented as a pattern of activation in a large collection of "units," each corresponding to a feature of the input item (a "distributed representation"). An output is represented as a pattern of activation on a second collection of units. Each input unit is connected to each output unit by a link whose weight is modifiable by training; sometimes one or more layers of "hidden" units stand between the input and the output layers. Some theorists treat the units as very simple models of neurons and treat the networks as models of the brain's massively parallel circuitry.

David Rumelhart and James McClelland (1986) devised a pattern associator model of the acquisition of past tense forms that acquired hundreds of regular and irregular verbs, and it generalized properly to dozens of new verbs that it had not been trained on. More strikingly, the model appeared to go through a U-shaped developmental sequence, first producing irregular verb forms correctly and later overregularizing them. But the model had no explicit representation of words and rules, nor any distinction between regular and irregular systems. It was just a single pattern associator network consisting of an array of input units, an array of output units, and a matrix of modifiable weighted links between every input and every output. The verb stem was represented by turning on a subset of input nodes, each corresponding to a sound pattern in the stem. For example, one unit might represent "vowel between two consonants"; another might represent "stop consonant at the end of the word"; a third might represent "high vowel between a sibilant consonant and a nonsibilant consonant." To represent the verb stem *pass,* the first and second units (and dozens of others) would be turned on; the third (and dozens of others) would be left off. This sent a signal across each of the links to the

output nodes, which represent the sounds of the past tense form. Each output node would sum its incoming signals and turn on if the sum exceeded a threshold; the output form would be the word most compatible with the set of active output nodes. During the learning phase the past tense form computed by the network is juxtaposed with the correct version provided by a "teacher," and the strengths of the links and thresholds are adjusted so as to reduce the difference. By recording and superimposing associations between stem sounds and past sounds, the model improves its performance and can generalize to new forms to the extent that their sounds overlap with old ones. This process is qualitatively the same for regular and irregular verbs: *stopped* is produced because input *op* units were linked to output *opped* units by previous verbs; *clung* is produced because *ing* was linked to *ung*.

Here is how Rumelhart and McClelland got their model to make overregularization errors after a period of correct performance, as we see in children. They reasoned that if children acquire verbs in order of decreasing frequency, they will develop a vocabulary with an increasing proportion of regular verbs as they begin to run out of the high-frequency irregulars and encounter more and more regular verbs. In particular, they noted that children seem to undergo an explosion in vocabulary development early in life, which could result in a sudden influx of a large number of regular verbs. They imagined an extreme case: The child first learns the dozen or so commonest verbs (which are almost all irregular), followed by an explosive acquisition of the next four hundred or so verbs, of which about 80 percent are regular. The network, bombarded by verbs showing the regular pattern, would strengthen many links between stem units and the units defining the sounds of the *-ed* ending. These newly modified links overwhelmed the existing links between features of irregular stems and the features of their idiosyncratic pasts, resulting in overregularization.

5.5.1 Testing Rules Versus Analogies

Is there any way to tell whether children's overregularization errors are due to overapplication of a mental rule or the effect of automatically analogizing from the patterns found in regular verbs? Marcus, Ullman, and I thought up several empirical tests.

First, we looked at the key assumption that during the child's development, overregularization was triggered by a sudden influx of regular verbs. What would this mean, concretely? It could mean that parents were suddenly using many more regular verbs in the past tense when talking to their children. We checked this assumption by counting the number of sentences with regular verbs and with irregular verbs used by the parents of Adam, Eve, and Sarah as the children grew. We found that the propor-

tion remained constant from 2 years of age to 5 years of age, about 30 percent regular. Thus a key assumption of the model turns out not to be true. This might seem odd, given that early conversations would presumably favor common irregular verbs like *make* and *do*, and later ones would make their way down to less common regular ones like *abhor* and *abrogate*, to take an extreme example. But *make*, *do*, *hold*, and so on are indispensable general-purpose verbs that people of all ages need to use; *abhor*, *abrogate*, and the like complete with *each other* for air time in conversation, leaving the proportion of regular verb forms more or less constant.

Even when we looked at the number of different kinds of verbs in the child's vocabulary (the "types"), as opposed to how often each verb is used (the "tokens"), we found no synchrony between the acquisition of regular verbs and the beginning of overregularization. There is an increase in the rate of vocabulary growth, to be sure, and that increase, as one would expect, results in an influx of regular verbs because most of the irregulars were learned at the beginning of development. But the spurt occurs when they are in their mid-to-late 1s, about a year earlier than the first overregularization error, which tends to occur in the mid-to-late 2s. During the ages in which children overregularize, the regulars are pouring in at a lower rate than when the children are using the irregulars correctly, exactly the opposite to what would have to happen to get a pattern-associator network to overregularize.

Second, we looked at the similarity between the overregularization errors and existing regular verbs. If *holded* is an analogy from similar verbs like *folded*, then the greater the number of similar regular verbs there are, and the more frequent each one is, the more likely the error should be, compared with other verbs. *Holded* might be a more common error than *drinked*, for example, because *holded* is attracted to frequent *folded* and to a lesser extent *scolded* and *moulded*, whereas *drinked* is attracted to relatively low-frequency *blinked* and nothing else. But when we did the calculations (correlating overregularization rates of the different verbs against the sums of the frequencies of the regular verbs that rhymed with them), we found no consistent correlation.

Third, we looked to see whether children's overregularizations simply followed similarity of sound, or whether they were made within the context of an entire system of mental grammar. The main idea was that if children go just by sound, they should overregularize a given verb sound at the same rate regardless of the grammatical status of the verb. But there are two cases where a given verb sound can behave in very different ways in terms of how children form their past tenses.

One of them, studied by Stromswold (1990), involves auxiliaries. As indicated in chapters 2 and 10, auxiliaries are a special kind of verb: They belong to the closed functional categories, as opposed to the open

lexical categories, and are completely sequestered from the kinds of grammatical rules that apply to verbs in general. (For example, you can't say *to must, musted, musting, have musted,* and so on). Conveniently, some of the auxiliaries in English also come in ordinary, nonauxiliary versions. *Do* can be an auxiliary, as in *I didn't eat* and *Did he eat,* but it can also be an ordinary verb meaning "act," as in *Do something!* and *We did it. Have* can be an auxiliary, as in *They have gone,* but it can also be an ordinary verb meaning "possess," as in *We have a Volvo.* And *be* can be an auxiliary, as in *He is running* or *It is being washed,* but it can also be an ordinary verb meaning "in the state of," as is *She is sick* and *He is a fool.* Interestingly, these verbs come in the same irregular forms whether they are being used as auxiliaries or as ordinary verbs: *does* and *did, has* and *had, am, is, are, was, were, been.* Moreover, the semantic relations are the same in all cases: the relation between *have a book* and *had a book* is the same as the relation between *have eaten* and *had eaten; I am tired* is to *I was tired* as *I am resting* is to *I was resting.* Clearly, there are too many of these parallelisms to be coincidental, and a parsimonious assumption is that the irregular forms of the main verb and of the auxiliary versions are stored in the same mechanism. But their susceptibility to overregularization is qualitatively different: In a sample of forty thousand child sentences containing these verbs, Stromswold found that the ordinary verb versions are overregularized at rates comparable to those I have been talking about (around 5 percent), whereas the auxiliary versions of the same verbs were *never* overregularized —0 errors. That is, children said *Look what he doed* but never *What doed he make;* they said *We haved a house* but never *We haved gone;* they said *He be'd sleepy* but never *He be'd sleeping.* This shows that they cannot just be analogizing on the basis of sound; if so, they would overregularize each of these sounds equally often, whether used as auxiliaries or as ordinary verbs.

Here is another interesting case where children care about a verb's grammar, not just its sound, in deciding how to put it in the past tense. An interesting quirk of grammar is that verbs intuitively perceived as derived from nouns or adjectives are always regular, even if they are similar or identical to an irregular verb. Thus one says *grandstanded,* not *grandstood; flied out* in baseball (hit *a fly*), not *flew out; high-sticked* in hockey, not *high-stuck.* The explanation is that an irregular memorized form is a linkage between two word roots, the atomic sound-meaning pairings stored in the mental lexicon; it is not a link between two words or sound patterns directly. *High-stuck* sounds silly because the verb is tacitly perceived as being based on the noun root (*hockey*) *stick,* and noun roots cannot be listed in the lexicon as having any past tense form (the past tense of a noun makes no sense semantically), let alone an irregular one. Because its root is not the verb *stick,* there is no route by which *stuck* can be made

available; to obtain a past tense form, the speaker must apply the regular rule, which serves as the default.

Now the question is, Are children sensitive to this subtle grammatical distinction? We found that they were (Kim, Marcus, Pinker, Hollander, and Coppola, 1994). We did a *wug*-test with a twist: half of the new verbs were homophonous with irregulars but were obviously based on nouns, like *to fly* in the sense of "to cover a piece of paper with flies" or *to ring* in the sense of "to put a ring on something." The 4-year-old children we tested regularized the verbs much more often when they were based on nouns (just like the adults who first said *grandstanded* and *flied out*) than when they were just plain irregular verbs, like "flying down the road" or "ringing a bell." Once again, children are not just overregularizing according to similarity of sound; they apply the regular suffix most often when the grammar of English demands it.

But we did find one correlation that *was* predicted by the connectionist pattern associator model. Though irregular verbs do not seem to be analogized to regulars, they might be analogized to *other irregulars*. We already know that irregular verbs fall into families of similar verbs: *ring-rang, drink-drank, sink-sank, shrink-shrank,* and so on; *grow-grew, know-knew, throw-threw, fly-flew,* and so on; *feed-fed, breed-bred, creep-crept,* and so on. Perhaps the irregulars come in these clusters because each member of the cluster serves as an analogy that helps in the memorization of all the others; it is easy to memorize *sink-sank* because it rhymes with *drink-drank*. (This is exactly what happens in connectionist pattern-associators because *sink* is represented by turning on many of the same units that are used by *drink*; thus, *sink* can exploit the connections to the irregular units for *ank* that have already been learned for *drank*). We found some good evidence for this effect when we looked at children's overregularization errors: The greater the number of rhyming irregular family members a verb had and the more common those rhyming members were in parents' speech, the less often the children overregularized them.

The analogizing of an irregular pattern, in fact, helps people not only to memorize verbs but also, once in a while, to apply it creatively to create new irregular forms. Children occasionally (about two-tenths of 1 percent of the opportunities) overapply irregular patterns, as in *wipe-wope* (compare *write-wrote*), *bring-brang* (compare *sing-sang*), and *write-writ* (compare *bite-bit*). And even adults are occasionally tempted, sometimes as a form of slang or humor (*squeeze-squoze,* or e.g., on the T-shirt that used to be sold by a famous Boston restaurant: "I got scrod at Legal Seafood"). A few of these kinds of innovations can make inroads into the language—like the form *snuck* I used earlier, a fairly recent creation (late nineteenth century) based on *stick-stuck*.

Does the demonstration of analogizing among irregular verbs contradict all those earlier findings showing that children were not engaging in connectionist-style analogizing? Not at all. Connectionist pattern associators may not be accurate models of the regular rule, but they may be accurate models of the memory system in which the irregulars are stored. Just as we noted early on that frequently heard words are easier to memorize, it seems that words that fall into mutually similar families are easier to memorize—and pattern associators display both those properties. Indeed, this might resolve an ongoing debate in cognitive science over whether rule systems or connectionist networks are better models of language processes. Both approaches might be right, but for different parts of the mind: rules for the combinatorial system underlying grammar, networks for the memory underlying word roots.

5.6 Conclusions

Why obsess over babytalk like *We holded the rabbits*? Because it is a simple enough example to study in great scientific detail but representative of important general problems in understanding language and language acquisition. In particular, it has led us to make the following observations.

1. A language contains rules that can generate an infinite number of brand-new word-forms and sentences.
2. When learning a language, children have to generalize from a finite sample of parental speech to the infinite set of sentences that define the language as a whole. Since there are an infinite number of ways to do this but only one is correct, children must be innately guided to the correct solution by having some kinds of principles governing the design of human language built in.
3. We can catch children in the act of generalizing when they use one of the general rules of English to create a form that they could not have heard from their parents. Children must be generalizing such a rule when they apply it to irregular verbs, as in *holded*, or to novel verbs in experiments, as in *ricked*.
4. Children command not just rules but memorized words, like *held*; they use the memorized irregular forms both prior to, and simultaneously with, the overgeneralized, rule-created forms.
5. Children's simultaneous use of correct and incorrect forms poses the puzzle of how they unlearn the incorrect forms, given that the incorrect forms are expressive and useful, and are not reliably corrected by parents.
6. The puzzle can be solved if children command one of the basic design features of language: the "Blocking" principle, whereby a rule is prevented from applying if there is a grammatically equivalent

irregular form in the memorized mental dictionary. As long as they can remember an irregular, they can stop producing the overregularized version.

7. The course of language development in this area can be explained straightforwardly, as an interaction between the innate organization of language (rules, words, and the Blocking principle that relates them) and the child's experience with parental speech. Early on, children just memorize words (*held*), though not perfectly. Later they formulate the regular past tense rule "add *-ed*" from memorized regular pairs like *walk-walked*. Now equipped with the rule, whenever they fail to retrieve an irregular past form from memory, they can apply the rule, resulting in an overregularization error. As they hear the irregular more and more often, they remember it better and better, block the rule more and more reliably, and make the errors less and less often.

8. There is considerable evidence for this hypothesis, in the form of demonstrations that children's language is organized much like adults', just with noisier retrieval of irregulars. The errors are uncommon, occur sporadically over all ages and verbs, and occur more often for less frequent irregular verbs. Children judge correct irregulars as better-sounding than their own errors. Adults, when they are put in the same circumstances that children invariably are in—insufficient experience hearing irregular forms—do what children do, namely, apply the regular rule to those forms. This is true both now and over the entire history of the language, and the regularizing tendencies of children and adults over the centuries have shaped the frequency structure of the language as we find it today.

9. Analogy plays a clear role in language. Children, and adults, occasionally analogize the pattern in one irregular verb to a new irregular verb (*write-wrote* → *bite-bote*). They also find it easier to memorize irregular verbs when they are similar to other irregular verbs. This analogizing is a hallmark of connectionist or parallel distributed processing pattern associators; it suggests that human memory might be like a pattern associator. The regular suffix, in contrast, does not seem to be applied by analogy: Children do not apply it more readily at the ages at which they have recently learned a bunch of regular verbs, and they do not apply it more readily to the verbs that are more similar to existing regular verbs. Moreover, generalizations of the regular pattern can bypass the word-sound system altogether and are applied (by children and adults) to irregular-sounding verbs that lack verb roots (that is, verbs derived from nouns). This shows that the regular suffix applies to the category of "verbs" across the board, not just verbs with familiar sounds.

Suggestions for Further Reading

A general introduction to language can be found in my book *The Language Instinct*, 1994. Chapters 2 and 9 are about language acquisition in general, and chapter 5 is about words and word structure, including regular and irregular inflection. The logical problem of language acquisition is discussed in detail by Wexler and Culicover 1980 and by Pinker 1979, 1984, 1987, 1989.

The latest facts and previous literature on overregularization in children are summarized in a *Monograph of the Society for Research in Child Development* by Gary Marcus et al. 1992. Marchman and Bates 1993 present an alternative account of the developmental data, which elicited a reply in one of the chapters of Marcus et al. 1992. Stemberger 1993 has found an interesting additional contributor to children's errors: children overregularize more if the vowel of the stem is more psychologically "natural" than the vowel of the past tense form.

The data on children's spontaneous speech errors come from an analysis of portions of the Child Language Data Exchange System, a computer archive of transcripts of children's conversations and related language materials. It is described in MacWhinney and Snow 1985, 1990, and in MacWhinney 1995. Roger Brown's study of the language of three children is reported in Brown 1973.

The general theory that the English past tense captures the two main psychological systems underlying language—rules of grammar and a memorized lexicon—is laid out in Pinker 1991, in Pinker and Prince 1994, and in Marcus et al. 1995, together with a review of the supporting evidence. Some of that evidence is reported in individual journal articles: On *flied out* and the process of regularizing irregular-sounding verbs derived from nouns, see Kim et al. 1991 and Kim et al. 1994. On analogies from regular and irregular verbs, see Prasada and Pinker 1993. For an interesting comparison of English-speaking children with German-speaking children, see Clahsen et al. 1993, and Clahsen, Marcus, and Bartke 1993. Marcus (in press) looks at children's overregularization of nouns. Stromswold 1990 compared overregularization of auxiliary and nonauxiliary verbs. Marcus (1993) examined parental "negative evidence" in language acquisition, a question first examined by Brown and Hanlon 1970.

The connectionist approach to cognition is explained in a recent textbook by Quinlan (1992). It was originally explained and defended in detail, with many case studies, in the two-volume set *Parallel Distributed Processing*, edited by David Rumelhart and James McClelland (1986). Rumelhart and McClelland's model of the acquisition of the past tense is described in one of the chapters in volume 2 of that set. The PDP volumes inspired a vigorous debate among cognitive psychologists, beginning with the three critical essays in a special issue of the journal *Cognition* in 1988 (volume 28) that was reprinted as a book by Pinker and Mehler 1988. The paper by Fodor and Pylyshyn 1988 reprinted in that volume is considered the principal general critique of the connectionist approach, and it led to replies by Chater and Oaksford 1990 and others (see the Quinlan book). The paper by Pinker and Prince (1988) examined Rumelhart and McClelland's past tense model as a concrete illustration of the strengths and weaknesses of connectionism, and it led to alternative connectionist past tense models defended by MacWhinney and Leinbach 1993, Plunkett and Marchman 1991, 1993, Hare and Elman 1993, and Daugherty and Seidenberg 1992.

Problems

Here is a list of the verb stems that have irregular past tense forms in modern standard American and British English:

alight, arise, awake, be, bear, beat, become, befall, beget, begin, behold, bend, beset, beshit, bespeak, bid, bind, bite, bleed, blow, break, breed, bring, build, burn, burst, buy, cast, catch, choose, cling, come, cost, creep, cut, deal, dig, dive, do, draw, dream, drink, drive, dwell, eat, fall, feed, feel, fight, find, fit, flee, fling, fly, forbear, forbid, foretell, forget, forgive, forgo, forsake, forswear, freeze, get, give, go, grind, grow, hang, have, hear, heave, hide, hit, hold, hurt, keep, kneel, knit, know, lead, leap, learn, leave, lend, let, lie, light, lose, make, mean, meet, mislead, mistake, partake, plead, prove, put, quit, read, rend, rid, ride, ring, rise, run, say, see, seek, sell, send, set, sew, shake, shear, shed, shine, shit, shoot, show, shrink, shut, sing, sink, sit, slay, sleep, slide, sling, slink, slit, smell, smite, sneak, sow, speak, speed, spell, spend, spill, spin, spit, split, spoil, spread, spring, stand, steal, stick, sting, stink, strew, stride, strike, string, strive, swear, sweep, swell, swim, swing, take, teach, tear, tell, think, thrive, throw, thrust, tread, undergo, understand, upset, wake, wear, weave, wed, weep, wet, win, wind, withdraw, withstand, wring, write

5.1 (a) Pull out a family of verbs (with five to twenty different verbs) such that the verb stems are all similar to one another and the past tense forms are all similar to one another. (b) Formulate a rule that would precisely convert the stems to the past tense forms for as many verbs as possible in the family (that is, a set of instructions of the form "if a verb has the following sequence of vowels and consonants, then convert such-and-such a vowel to such-and-such a vowel, and such-and-such a consonant to such-and-such a consonant"). There will be positive exceptions of your rule—verbs that fit the "if" part but whose past tense is different from what you get when you apply the "then" part. (c) Give an example of a verb that is a positive exception to your rule, and show how it is one. There will also be negative exceptions: verbs that don't fit the "if" part of your rule, but which have a past tense form that is similar to the one that would have been created by the "then" part. (e) Give an example of a verb that is a negative exception to your rule, and show how it is one. (f) Given what you now know about this family of verbs, what can you say about how a child might learn their past tense forms?

5.2 (a) List all the verbs for which you would use a *regular* past tense form (for example, most people would say *heaved* as the past tense for *heave*, but it is listed here because the form *hove* exists in English as well). List their irregular forms as well (use a dictionary if you have to). (b) Now list an equal number of verbs for which you would *never* use the regular form (for example, no adult would say *telled*). (c) Are there any systematic differences between the (a) verbs and the (b) verbs? (d) Could those differences explain why a child might overregularize any of the verbs in your (a) list?

5.3 In English, one can turn a noun into a verb in several ways: *I pocketed the money* (put it into a pocket); *She peeled the carrot* (took the peel off); *He elbowed his way in* (used his elbow as an instrument); *I'm all gamed out* (I've had enough games); *She chickened out* (acted like a chicken). (a) Find three verbs in the list of irregular verbs that are homophonous (have the same sound) as a noun. (b) Find, or make up, a verb based on that noun. (c) Make up a sentence in which the verb is used in the past tense. (d) Is the past tense form the same as, or different from, that of the original irregular verb you chose? (e) If there is a difference, explain it.

References

Berko, J. (1958). The child's learning of English morphology. *Word* 14, 150–177.

Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.

Brown, R., and C. Hanlon (1970). Derivational complexity and order of acquisition in child speech. In J. R. Hayes, ed., *Cognition and the development of language*. New York: Wiley.

Chater, N., and M. Oaksford (1990). Autonomy, implementation and cognitive architecture: A reply to Fodor and Pylyshyn. *Cognition* 34, 93–108.

Clahsen, H., M. Rothweiler, A. Woest, and G. F. Marcus (1993). Regular and irregular inflection in the acquisition of German noun plurals. *Cognition* 45, 225–255.

Clahsen, H., G. Marcus, and S. Bartke (1993). Compounding and inflection in German child language. *Essex Research Reports in Linguistics* 1, Colchester, England.

Daugherty, K., and M. Seidenberg (1992). Rules or connections? The past tense revisited. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 259–264. Hillsdale, NJ: Erlbaum.

Fodor, J. A., and Z. Pylyshyn (1988). Connectionism and cognitive architecture. A critical analysis. *Cognition* 28, 3–71.

Hare, M., and J. Elman (1993). A connectionist account of English inflectional morphology: Evidence from language change. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 265–270. Hillsdale, NJ: Erlbaum.

Kim, J. J., S. Pinker, A. S. Prince, and S. Prasada (1991). Why no mere mortal has ever flown out to center field. *Cognitive Science* 15, 173–218.

Kim, J. J., G. Marcus, S. Pinker, M. Hollander, and M. Coppola (1994). Sensitivity of children's inflection to morphological structure. *Journal of Child Language* 21, 173–209.

Kuczaj, S. A. (1978). Children's judgments of grammatical and ungrammatical past-tense forms. *Child Development* 49, 319–326.

MacWhinney, B. (1995). *The CHILDES Project: Computational tools for analyzing talk* (2nd ed.). Hillsdale, NJ: Erlbaum.

MacWhinney, B., and J. Leinbach (1991). Implementations are not conceptualizations: Revising the verb learning model. *Cognition* 40, 121–157.

MacWhinney, B., and C. E. Snow (1985). The Child Language Data Exchange System. *Journal of Child Language* 12, 271–296.

MacWhinney, B., and C. E. Snow (1990). The Child Language Data Exchange System: An update. *Journal of Child Language* 17, 457–472.

Marchman, V., and E. Bates (1993). Continuity in lexical and morphological development: A test of the critical mass hypothesis. *Journal of Child Language* 21, 339–366.

Marcus, G. F. (1993). Negative evidence in language acquisition. *Cognition* 46, 53–85.

Marcus, G. F. (in press). Children's overregularization of English plurals: A quantitative analysis.

Marcus, G. F., U. Brinkmann, H. Clahsen, R. Wiese, and S. Pinker (1995). German inflection: The exception that proves the rule. Submitted to *Cognitive Psychology*.

Marcus, G., S. Pinker, M. Ullman, M. Hollander, T. J. Rosen, and F. Xu (1992). Overregularization in language acquisition. *Monographs of the Society for Research in Child Development* 57 (4, Serial No. 228).

Pinker, S. (1979). Formal models of language learning. *Cognition* 1, 217–283.

Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.

Pinker, S. (1987). The bootstrapping problem in language acquisition. In B. MacWhinney, ed., *Mechanisms of language acquisition*. Hillsdale, NJ: Erlbaum.

Pinker, S. (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.

Pinker, S. (1991). Rules of language. *Science* 253, 530–535. Reprinted in P. Bloom, ed., *Language acquisition*: Core readings. Cambridge, MA: MIT Press.

Pinker, S. (1994a). *The language instinct*. New York: Morrow. London: Penguin.

Pinker, S., and J. Mehler, eds. (1988). *Connections and symbols*. Cambridge, MA: MIT Press.

Pinker, S., and A. Prince (1988). On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* 28, 73–193.

Pinker, S., and A. Prince (1994). Regular and irregular morphology and the psychological status of rules of grammar. In S. D. Lima, R. L. Corrigan, and G. K. Iverson (eds.), *The reality of linguistic rules*. Philadelphia: Benjamin.

Plunkett, K., and V. Marchman (1991). U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition* 38, 43–102.

Plunkett, K., and V. Marchman (1993). From rote learning to system building. *Cognition* 48, 21–69.

Prasada, S., and S. Pinker (1993). Generalizations of regular and irregular morphological patterns. *Language and Cognitive Processes* 8, 1–56.

Quinlan, P. (1992). *An introduction to connectionist modeling*. Hillsdale, NJ: Erlbaum.

Rumelhart, D., and J. McClelland (1986). On learning the past tenses of English verbs. Implicit rules or parallel distributed processing? In J. McClelland, D. Rumelhart, and the PDP Research Group, *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Rumelhart, D., J. McClelland, and the PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.

Stemberger, J. (1993). Vowel dominance in overregularizations. *Journal of Child Language* 20, 503–521.

Stromswold, K. J. (1990). Learnability and the acquisition of auxiliaries. Doctoral dissertation, Department of Brain and Cognitive Sciences, MIT.

Stromswold, K. J. (1995). What a mute child tells us about language. Unpublished manuscript, Rutgers University.

Ullman, M. (1993). The computation and neural localization of inflectional morphology. Doctoral dissertation, Department of Brain and Cognitive Sciences, MIT.

Ullman, M., S. Corkin, S. Pinker, M. Coppola, J. Locascio, and J. H. Growdon (1993). Neural modularity in language: Evidence from Alzheimer's and Parkinson's disease. (abstract) Paper presented at the 23rd Annual Meeting of the Society for Neuroscience, Washington DC.

Wexler, K., and P. Culicover (1980). *Formal principles of language acquisition*. Cambridge, MA: MIT Press.